arXiv:1612.04844v1 [cs.CV] 14 Dec 2016

The More You Know: Using Knowledge Graphs for Image Classification

Kenneth Marino, Ruslan Salakhutdinov, Abhinav Gupta
Carnegie Mellon University

5000 Forbes Ave, Pittsburgh, PA 15213 {kdmarino, rsalakhu, abhinavg}@cs.cmu.edu

Abstract

Humans have the remarkable capability to learn a large variety of visual concepts, often with very few examples, whereas current state-of-the-art vision algorithms require hundreds or thousands of examples per category and struggle with ambiguity. One characteristic that sets humans apart is our ability to acquire knowledge about the world and reason using this knowledge. This paper investigates the use of structured prior knowledge in the form of knowledge graphs and shows that using this knowledge improves performance on image classification. Specifically, we introduce the Graph Search Neural Network as a way of efficiently incorporating large knowledge graphs into a fully end-to-end learning system. We show in a number of experiments that our method outperforms baselines for multi-label classification, even under low data and few-shot settings.

1. Introduction

Our world contains millions of visual concepts understood by humans. These often are ambiguous (tomatoes can be red or green), overlap (vehicles includes both cars and planes) and have dozens or hundreds of subcategories (thousands of specific kinds of insects). While some visual concepts are very common such as person or car, most categories have many fewer examples, forming a long-tail distribution [39]. And yet, even when only shown a few or even one example, humans have the remarkable ability to recognize these categories with high accuracy. In contrast, while modern learning-based approaches can recognize some categories with high accuracy, it usually requires hundreds or thousands of labeled examples for each of these categories. Given how large, complex and dynamic the space of visual concepts is, this approach of building large datasets for every concept is unscalable. Therefore, we need to ask what humans have that current approaches do not.

One possible answer to this is structured knowledge and reasoning. Humans are not merely appearance-based classifiers; we gain knowledge of the world from experience and language. We use this knowledge in our everyday lives to recognize unfamiliar objects. For instance, we might have read in a book about the "elephant shrew" (maybe even seen an example) and will have gained knowledge that is useful for recognizing one. Figure 1 illustrates how we might use our knowledge about the world in this problem. We might know that an elephant shrew looks like a mouse, has a trunk and a tail, is native to Africa, and is often found in bushes. With this information, we could probably identify the elephant shrew if we saw one in the wild. We do this by first recognizing (we see a small mouse-like object with a trunk in a bush), recalling knowledge (we think of animals we have heard of and their parts, habitat, and characteristics) and then reasoning (it is an elephant shrew because it has a trunk and a tail, and looks like a mouse while mice and elephant do not have all these characteristics). With this information, even if we have only seen one or two pictures of this animal, we would easily be able to classify it.

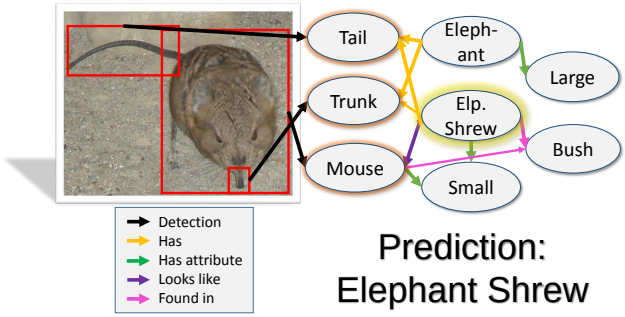

Figure 1. Example of how semantic knowledge about the world aids classification. Here we see an elephant shrew. Humans are able to make the correct classification based on what we know about the elephant shrew and other similar animals.

The focus of this work is to exploit structured visual knowledge and reasoning to improve image classification. Recently, there has been a lot of work that has focused on building knowledge graphs that might be useful for images problems. For example, Never Ending Image Learner [38]

learns relationships between visual concepts directly from images from the web. Similarly, Never Ending Language Learner [4] learns relationships between semantic categories, but in this case by reading text on the web. There are also a number of human labeled knowledge bases such as WordNet [23], which was the starting point for the ImageNet [28] categories, and Visual Genome [15] which has human annotated scene graphs for each image. However, given the scale and ambiguity issues, these knowledge graphs are quite noisy. The knowledge exists; the question is how to use it effectively in the presence of noise.

There has been a lot of work in end-to-end learning on graphs or neural network trained on graphs [30, 3, 6, 11, 24, 22, 9, 21]. Most of these approaches either extract features from the graph or they learn a propagation model that transfers evidence between nodes conditional on the type of relationship. An example of this is the Gated Graph Neural Network [18] which takes an arbitrary graph as input. Given some initialization (annotation specific to the task such as starting and ending node for shortest path), it learns how to propagate information and predict the output for every node in the graph. This approach has been shown to solve basic logical tasks as well as program verification.

Our work improves on this model and adapts end-to-end graph neural networks to the image classification task. We introduce the Graph Search Neural Network (GSNN) which uses features from the image to efficiently annotate the graph, select a relevant subset of the input graph and predict outputs on nodes representing visual concepts. These output states are then used to classify the image. GSNN learns a propagation model which reasons about different types of relationships and concepts to produce outputs on the nodes which are then used for image classification. Our new architecture mitigates the computational issues with the Gated Graph Neural Networks for large graphs which allows our model to be efficiently trained for image tasks using large knowledge graphs. We show how our GSNN model is effective at reasoning about concepts to improve image classification tasks across the entire spectrum—from one-shot learning to large-scale dataset learning. Importantly, our GSNN model is also able to provide explanations on classifications by following how the information is propagated in the graph.

The major contributions of this work are (a) the introduction of the GSNN as a way of incorporating potentially large knowledge graphs into an end-to-end learning system that is computationally feasible for large graphs; (b) a framework for using noisy knowledge graphs for image classification; (c) the ability to explain our image classifications by using the propagation model; (d) the introduction of a new subset of Visual Genome designed to test 1-shot and few-shot learning without any overlap with classes from ImageNet [28] or COCO [19]. Our method significantly outperforms baselines for multi-label classification both in full data and low-data settings.

2. Related Work

Learning knowledge graphs [38, 4, 29] and using graphs for visual reasoning [39, 20] has recently been of interest to the vision community. For reasoning on graphs, several approaches have been studied. For example, [40] collects a knowledge base and then queries this knowledge base to do first-order probabilistic reasoning to predict affordances. [20] builds a graph of exemplars for different categories and uses the spatial relationships to perform contextual reasoning. Approaches such as [17] use random walks on the graphs to learn patterns of edges while performing the walk and predict new edges in the knowledge graph. There has also been some work using a knowledge base for image retrieval [12] or answering visual queries [41], but these works are focused on building and then querying knowledge bases rather than using existing knowledge bases as side information for some vision task.

However, none of these approaches have been learned in an end-to-end manner and the propagation model on the graph is mostly hand-crafted. More recently, learning from knowledge graphs using neural networks and other end-to-end learning systems to perform reasoning has become an active area of research. Several works treat graphs as a special case of a convolutional input where, instead of pixel inputs connected to pixels in a grid, we define the inputs as connected by an input graph, relying on either some global graph structure or doings some sort of pre-processing on graph edges [3, 6, 11, 24]. However, most of these approaches have been tried on smaller graphs such as molecular datasets. In vision problems problems, these graphs encode contextual and common-sense relationships and are significantly larger, leading to scalability issues.

Li and Zemel present Graph Gated Neural Networks (GGNN) [18] which uses neural networks on graph structured data. This paper (itself an extension of Graph Neural Networks [30]) serves as the foundation for our Graph Search Neural Network (GSNN). Several papers have found success using variants of Graph Neural Networks applied to various simple domains such as quantitative structure-property relationship (QSPR) analysis in chemistry [22] and subgraph matching and other graph problems on toy datasets [9]. GGNN is a fully end-to-end network that takes as input a directed graph and outputs either a classification over the entire graph or an output for each node. For instance, for the problem of graph reachability, GGNN is given a graph, a start node and end node, and the GGNN will have to output whether the end node is reachable from the start node. They show results for simple logical tasks on graphs and more complex tasks such as program verification.

There is also a substantial amount of work on various types of kernels defined for graphs [35] such as diffusion kernels [14], graphlet kernels [32], Weisfeiler-Lehman graph kernels [31], deep graph kernels [26], graph invariant kernels [25] and shortest-path kernels [1]. The methods have various ways of exploiting common graph structures, however, these approaches are only helpful for kernel-based approaches such as SVMs which do not compare well with neural network architectures in vision.

Our work is also related to attribute approaches [8] to vision such as [16] which uses a fixed set of binary attributes to do zero-shot prediction, [33] which uses attributes shared across categories to prevent semantic drift in semi-supervised learning and [5] which automatically discovers attributes and uses them for fine-grained classification. Our work also uses attribute relationships that appear in our knowledge graphs, but also uses relationships between objects and reasons directly on graphs rather than using object-attribute pairs directly.

We further evaluate our model on one and few-shot learning. There have been a few works in this area recently such as [2] which inverts the process that creates the input and [37] which learns new categories by finding common patches between classes. Both evaluate on handwritten characters and it is unclear how well they would scale to realistic images. [36] improves few-shot classification by training a network on a high-data image dataset and then learning a model-to-model transformation onto the new dataset.

3. Methodology

3.1. Graph Gated Neural Networks

The idea of GGNN is that given a graph with N nodes, we want to produce some output which can either be an output for every graph node $o_1, o_2, ...o_N$ or a global output o_G. This is done by learning a propagation model similar to an LSTM. For each node in the graph v, we have a hidden state representation $h_v^{(t)}$ at every time step t. We start at $t = 0$ with initial hidden states x_v that depends on the problem. For instance, for learning graph reachability, this might be a two bit vector that indicates whether a node is the source or destination node. In case of visual knowledge graph reasoning, x_v can be a one bit activation representing the confidence of a category being present based on an object detector or classifier.

Next, we use the structure of our graph, encoded in a matrix A which serves to retrieve the hidden states of adjacent nodes based on the edge types between them. The hidden states are then updated by a gated update module similar to an LSTM. The basic recurrence for this **propagation network** is

$$h_v^{(1)} = [x_v^T, 0]^T \tag{1}$$

$$a_v^{(t)} = A_v^T [h_1^{(t-1)} ... h_N^{(t-1)}]^T + b \tag{2}$$

$$z_v^t = \sigma(W^z a_v^{(t)} + U^z h_v^{(t-1)}) \tag{3}$$

$$r_v^t = \sigma(W^r a_v^{(t)} + U^r h_v^{(t-1)}) \tag{4}$$

$$\widetilde{h_v^t} = tanh(W a_v^{(t)} + U(r_v^t \odot h_v^{(t-1)})) \tag{5}$$

$$h_v^{(t)} = (1 - z_v^t) \odot h_v^{(t-1)} + z_v^t \odot \widetilde{h_v^t} \tag{6}$$

where Eq 1 is the initialization of the hidden state with x_v and empty dimensions. Eq 2 shows the propagation updates from adjacent nodes. Eq (3-6) combine the information from adjacent nodes and current hidden state of the nodes to compute the next hidden state.

After T time steps, we have our final hidden states. The node level outputs can then just be computed as

$$o_v = g(h_v^{(T)}, x_v) \tag{7}$$

where g is a fully connected network, the **output network**, and x_v is the original annotation for the node.

3.2. Graph Search Neural Network

The biggest problem in adapting GGNN for image tasks is computational scalability. NEIL [38] for example has over 2000 concepts, and NELL [4] has over 2M confident beliefs. Even after pruning to our task, these graphs would still be huge. Forward propagation on the standard GGNN is $O(N^2)$ to the number of nodes N and backward propagation is $O(N^T)$ where T is the number of propagation steps. We perform simple experiments on GGNNs on synthetic graphs and find that after more than about 500 nodes, a forward and backward pass takes over 1 second on a single instance, even when making generous parameter assumptions. On 2,000 nodes, it takes well over a minute for a single image. Clearly using GGNN out of the box is infeasible.

Our solution to this problem is the Graph Search Neural Network (GSNN). As the name might imply, the idea is that rather than performing our recurrent update over all of the nodes of the graph at once, we start with some initial nodes based on our input and only choose to expand nodes which are useful for the final output. Thus, we only compute the update steps over a subset of the graph. So how do we select which subset of nodes to initialize the graph with? During training and testing, we determine initial nodes in the graph based on likelihood of the concept being present as determined by an object detector or classifier. For our experiments, we use Faster R-CNN [27] for each of the 80 COCO categories. For scores over some chosen threshold, we choose the corresponding nodes in the graph as our initial set of active nodes.

Once we have initial nodes, we also add the nodes adjacent to the initial nodes to the active set. Given our initial nodes, we want to first propagate the beliefs about our initial nodes to all of the adjacent nodes. After the first time step, however, we need a way of deciding which nodes to

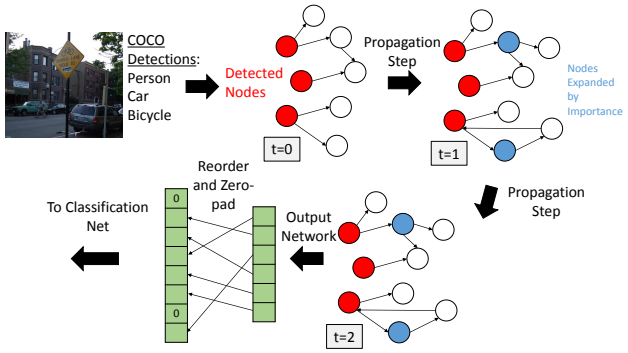

Figure 2. Graph Search Neural Network expansion. Starts with detected nodes and expands neighbors. Adds nodes adjacent to expand nodes predicted by importance net.

expand next. We therefore learn a per-node score that estimates how "important" that node is. After each propagation step, for every node in our current graph, we predict an importance score

$$i_v^{(t)} = g_i(h_v, x_v) \tag{8}$$

where g_i is a learned network, the **importance network**.

Once we have values of i_v, we take the top P scoring nodes that have never been expanded and add them to our expanded set, and add all nodes adjacent to those nodes to our active set. Figure 2 illustrates this expansion. At $t = 0$ only the detected nodes and their neighbors are expanded. At $t = 1$ we expand chose nodes based on importance values and add their neighbors to the graph. At the final time step ($t = 2$ in this case) we compute the per-node-output and re-order the outputs into the final classification net.

To train the importance net, we assign target importance value to each node in the graph for a given image. Nodes corresponding to ground-truth concepts in an image are assigned the important value 1. The neighbors of these nodes are assigned a value of γ. Nodes which are two-hop away have value γ^2 and so on. The idea is that nodes closest to the final output are the most important to expand.

We now have an end-to-end network which takes as input a set of initial nodes and annotations and outputs a per-node output for each of the active nodes in the graph. It consists of three sets of networks: the propagation net, the importance net, and the output net. The final loss from the image problem can be backpropagated from the final output of the pipeline back through the output net and the importance loss is backpropagated through each of the importance outputs. See Figure 3 to see the GSNN architecture.

One final detail is the addition of a "node bias" into GSNN. In GGNN, the per-node output function $g(h_v^{(T)}, x_v)$ takes in the hidden state and initial annotation of the node v to compute its output. In a certain sense it is agnostic to the meaning of the node. That is, at train or test time, GSNN takes in a graph it has perhaps never seen before, and some

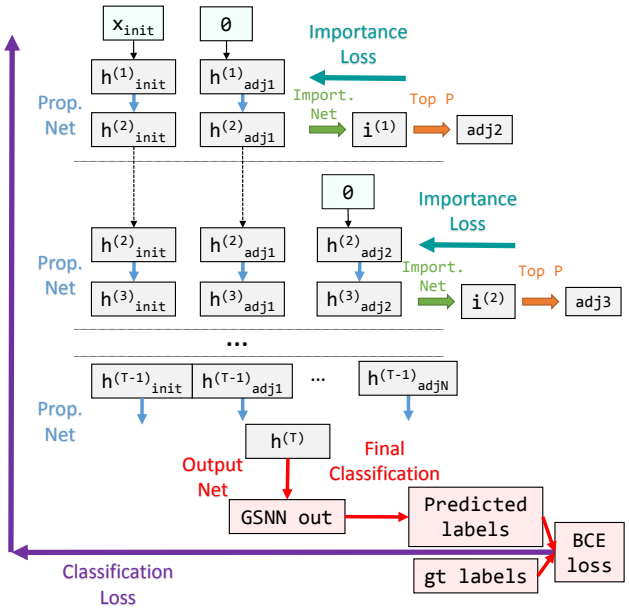

Figure 3. Graph Search Neural Network diagram. Shows initialization of hidden states, addition of new nodes as graph is expanded and the flow of losses through the output, propagation and importance nets.

initial annotations x_v for each node. It then uses the structure of the graph to propagate those annotations through the network and then compute an output. The nodes of the graph could have represented anything from human relationships to a computer program. However, in vision problems, the fact that a particular node represents "horse" or "cat" will probably be relevant, and we can also constrain ourselves to a static graph over image concepts. Hence we introduce node bias terms that, for every node in our graph, has some learned values. Our output equations are now $g(h_v^{(T)}, x_v, n_v)$ where n_v is a bias term that is tied to a particular node in the overall graph. This value is stored in a table and its value are updated by backpropagation.

3.3. Image pipeline and baselines

Another problem we face adapting graph networks for vision problems is how to incorporate the graph network into an image pipeline. For classification at least, this is fairly straightforward. We simply take the output of the graph network, reorder it so that nodes always appear in the same order into the final network, and zero pad any nodes that were not expanded by our algorithm. Therefore, if we have a graph of 316 nodes and each node uses a 5-dim hidden variable. We create a 1580-dim feature vector from the graph. We also concatenate this feature vector with fc7 layer (4096-dim) of fine-tuned VGG-16 network [34] and top-score for each COCO category predicted by Faster R-CNN (80-dim). This 5756-dim feature vector is then fed into 1-layer final classification network which is

trained with dropout.

For baselines, we compare to: (1) VGG Baseline - feed just fc7 into final classification net; (2) Detection Baseline - feed fc7 and top COCO scores into final classification net.

4. Results

4.1. Datasets and Graphs

For our experiments, we wanted to test on a dataset that better represents the complex, noisy visual world with its many different kinds of objects, where labels are potentially ambiguous and overlapping, and categories fall into a long-tail distribution [39]. Humans do well in this setting, but vision algorithms still struggle with it, and this is where we would expect knowledge-based approaches to help the most. To this end, we chose the Visual Genome dataset [15] v1.0.

Visual Genome contains over 100,000 natural images from the Internet. Each image is labeled with objects, attributes and relationships between objects entered by human annotators. Annotators could enter any object in the image rather than from a predefined list, so as a result there are thousands of object labels with some being more common and most having many fewer examples. There are on average 21 labeled objects in an image, so compared to datasets such as ImageNet [28], COCO [19] or PASCAL [7], the scenes we are considering are far more complex.

To evaluate our method, we perform two different experiments. In the first experiment, we evaluate the performance on multi-label classification. In the second experiment, we explore the performance on low-shot recognition.

4.2. Multi-Label Classification

For the first experiment, we create a subset from Visual Genome which we call Visual Genome multi-label dataset or VGML. In VGML, we take the 200 most common objects in the dataset and the 100 most common attributes and also add any COCO categories not in those 300 for a total of 316 visual concepts. Our task is then multi-label classification: for each image predict which subset of the 316 total categories appear in the scene. We randomly split the images into a roughly 80-20 train/test split. Since we used pre-trained detectors from COCO, we ensure none of our test images are from COCO training images.

We also use Visual Genome as a source for our knowledge graph. Using only the train split, we build a knowledge graph connecting the concepts using the most common object-attribute and object-object relationships in the dataset. Specifically, we counted how often an object/object relationship or object/attribute pair occurred in the training set, and pruned any edges that had fewer than 200 instances. This leaves us with a graph over all of the images with each edge being a common relationship. The idea is that we would get very common relationships (such as grass is green or person wears clothes) but not relationships that are rare and only occur in single images (such as person rides zebra). See Appendix for more details.

The Visual Genome graphs are useful for our problem because they contain scene-level relationships between objects, e.g. person wears pants or fire hydrant is red and thus allow the graph network to reason about what is in a scene. However, it does not contain useful semantic relationships. For instance, it might be helpful to know that dog is an animal if our visual system sees a dog and one of our labels is animal. To address this, we also create a version of graph by fusing the Visual Genome Graphs with Wordnet [23]. Using the subset of Wordnet from [10], we first collect new nodes in WordNet not in our output label by including those which directly connect to our output labels and thus likely to be relevant and add them to a combined graph. We then take all of the WordNet edges between these nodes and add them to our combined graph (see Appendix for more details on these graphs).

4.3. Low-Shot Recognition

The second dataset split, which we call the Visual Genome few-shot multi-label dataset or VGFS, tries to make the problem even more difficult. We would like to have a multi-label dataset where we had only a few examples per class and classes that have no overlap to ImageNet (used for learning visual features) or COCO (used for activating initial nodes in the graph).

From Visual Genome, we select 100 classes that are distinct from any ImageNet or COCO classes (see Appendix for the complete list). For each class, we hand-select up to five training images per category for the dataset. For the "one-shot" experiments, we only have the 100 images corresponding to one image selected per category, and for the 5-shot, we have the 500 images. Note that because it is a multi-label dataset, we end up with more than one example per class since we choose 1 image per category, but the category may happen to appear in the image for another category as well. For the 1-shot dataset, 46 categories have one image, 22 have 2, 15 have 3, and 17 have 4 or more. We divide the remaining 53,155 images that have at least one of the 100 categories and are not in the COCO 2014 Train split into a test and holdout set. We use the holdout set to create a graph over the 100 new categories as well as the categories from VGML (see Appendix for more details on the dataset and the knowledge graph). This knowledge graph represents common relationships among categories and can be learned from other sources such as NEIL [38]. For our experiments, we also use another version of knowledge graph by fusing WordNet knowledge similar to VGML setting.

Method	500	1k	5k	full
VGG	10.28	12.09	18.58	19.15
VGG+Det	10.54	13.05	22.42	23.52
GSNN-VG	11.63	15.19	21.51	23.49
GSNN-VG+WN	**11.74**	**16.14**	**24.75**	**26.70**

Table 1. Mean Average Precision for multi-label classification on Visual Genome Multi-Label dataset using 500, 1,000, 5,000 training examples and the full training dataset. Numbers for VGG baseline, VGG baseline with detections, GSNN using Visual Genome graph and GSNN using combined Visual Genome and WordNet graph.

Method	1-shot	5-shot
VGG	5.96	8.07
VGG+Det	4.77	9.09
GSNN-VG	6.60	9.85
GSNN-VG+WN	**7.30**	**11.11**

Table 2. Mean Average Precision for multi-label classification on Visual Genome few-shot dataset. Results shown for 1-shot (100 total images) and 5-shot (500 total images).

4.4. Quantitative Evaluation

We first report results on the VGML dataset. We train the models using ADAM [13] with an initial learning rate of 10^{-3} for all networks, except the pre-trained VGG where we use an initial learning rate of 10^{-4}, and an initial momentum of 0.9, except the GSNN which used 0. We set our GSNN hidden state size to 10, importance discount factor γ to 0.3, number of steps T to 2, initial confidence threshold to 0.1 and our expand number P to 5. Our GSNN importance and output networks are single layer networks with sigmoid activations. All networks were trained for 2,653 iterations with a batch size of 32.

Table 1 shows the result of our method on the multi-label classification task. We show the performance for different training dataset sizes. In all experiments, we see that the GSNN model with the combined graph outperforms all baselines by a significant margin. In the low data regime, both GSNN nets outperform baselines, likely because the baseline models cannot learn the relationships between all of the output categories, as they have not seen enough examples of many of the categories. In the full data regime, baselines achieve comparable performance to the GSNN with just the Visual Genome graph but are outperformed by the combined graph. This suggests that including the outside semantic knowledge from WordNet and performing explicit reasoning on a knowledge graph allows our model to learn better representations compared to the other models.

Table 2 shows the result of our method on the few shot task (VGFS). We use the same learning and network parameters as with our previous experiments for 20 epochs. We choose our GSNN hidden state size again as 10, our expand number P as 5 for 1-shot and 20 for 5-shot, our importance discount factor γ as 0.3, our initial confidence threshold as 0.1 and our number of steps T as 2. The P for 1-shot is chosen lower to avoid overfitting since we had many fewer examples to train with.

The dataset is much smaller and the categories are not in COCO or ImageNet, so all models achieve lower precision. Nevertheless, our GSNN models show improvements over baselines, exploiting the information from our knowledge graphs, especially with the combined graph.

As one might suspect, our method does not perform uniformly on all categories, but rather does better on some categories and worse on others. Figure 4 shows the differences in average precision for each category between our GSNN model with the combined graph and the detection baseline for the 1k VGML experiment. Figure 5 shows the same for our 5-shot VGFS experiment. To see this analysis for the other experiments, see Appendix. As expected, since the GSNN models have higher mAP, for the majority of categories, the AP improves. Performance on some classes improves greatly, such as "laptop" in our 1k VGML experiment and "runway" in our VGFS 5-shot experiment. Others however, such as "field" in our VGML 1k experiment perform worse. In the next section, we analyze our GSNN models on several examples to try to gain a better intuition as to what the GSNN model is doing and why it does well or poorly on certain examples.

4.5. Qualitative Evaluation

One way to analyse the GSNN is to look at the sensitivities of parameters in our model with respect to a particular output. Given a single image I, and a single label of interest y_i that appears in the image, we would like to know how information travels through the GSNN and what nodes and edges it uses. We examined the sensitivity of the output to hidden states and detections by computing the partial derivatives $\frac{\partial y_i}{\partial h^{(1)}}$ $\frac{\partial y_i}{\partial h^{(2)}}$ $\frac{\partial y_i}{\partial x_{det}}$ with respect to the category of interest. These values tell us how a small change in the hidden state of a particular node affects a particular output. We would expect to see, for instance, that for labeling elephant, we see a high sensitivity for the hidden states corresponding to grey and trunk.

In this section, we show the sensitivity analysis for the GSNN combined graph model on the VGML 1k experiment and the VGFS 5-shot experiments. In particular, we examine classes that performed well under GSNN and a few that performed poorly to try to get a better intuition into why some categories improve more over baselines. More examples for the other experiments are provided in the Appendix.

Figure 6 shows the graph sensitivity analysis for the experiments with VGML experiment on the left and VGFS on the right, showing four successful detections by the GSNN network and two failures. Each example shows the image,

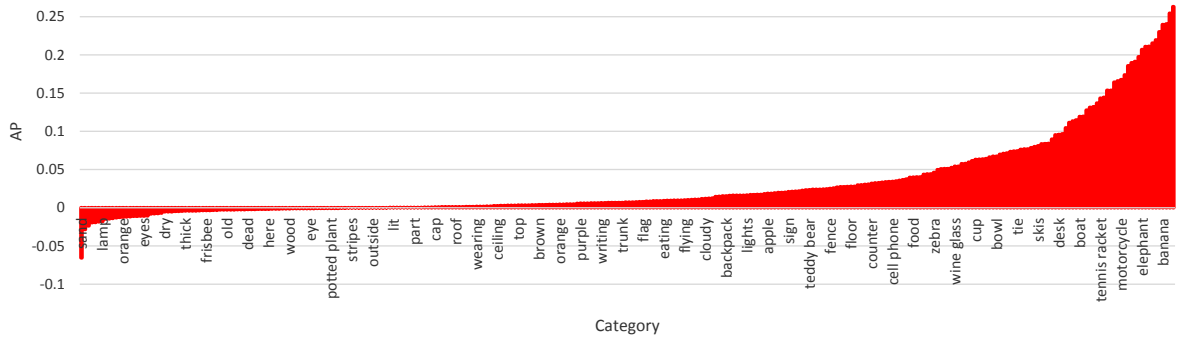

Figure 4. Difference in Average Precision for each of the 316 labels in VGML between our GSNN combined graph model and detection baseline for 1k experiment. Top categories: broccoli, laptop, train, banana, plane. Bottom categories: sand, tree, snow, beach, field.

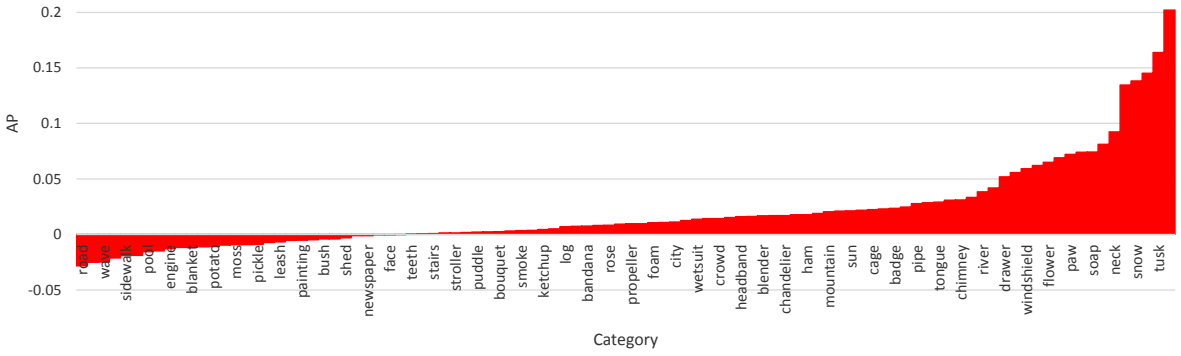

Figure 5. Difference in Average Precision for each of the 100 labels in VGFS between our GSNN combined graph model and detection baseline for 5-shot experiment. Top categories: beak, tusk, runway, snow, headboard. Bottom categories: road, onion, wave, bacon, sidewalk.

the ground truth output we are analyzing and the sensitivities of the concept of interest with respect to the hidden states of the graph or detections. For convenience, we display the names of the top detections or hidden states. We also show part of the graph that was expanded, to see what relationships GSNN was using.

For the VGML experiment, the top left of Figure 6 shows that using keyboard, laptop, cat and dining table detections, GSNN is able to reason that because laptop and keyboard occur together that it is looking at a desk, making dining table and teddy bear less likely. In the next hidden state the attributes for laptop, silver and open become more important. The middle left shows that the airplane detection is important. The GSNN in particular uses the "has" connection between airplane and wing. The bottom left shows the failure case for the 1k. It uses the person detection very strongly and in the next hidden states looks at less helpful parts of the graph, such as head and building that are not related to our output class "field." On the top right for the VGFS 5-shot experiment, we see that our network highly correlates the airplane detection, and reasons that runway, smoke and sky are also in the figure using the graph. On the middle right we see that for neck, it uses the giraffe detection and uses the fact that there is a "has" edge in the graph from giraffes to neck. It also weakly used the couch detection, but the graph reinforces the stronger giraffe detection during expansion. The bottom right shows a missclassified 5-shot example. While the GSNN is sensitive to the pizza detection which has a connection to onion, it is also fooled by a high vase detection sensitivity and expanded to unrelated nodes bouquet and flower.

5. Conclusion

In this paper, we present the Graph Search Neural Network (GSNN) as a way of efficiently using knowledge graphs as extra information to improve image classification. We show that even in low-data settings, our model performs well. We provide analysis that examines the flow of information through the GSNN and provides insights into why our model improves performance. We hope that this work provides a step towards bringing symbolic reasoning into traditional feed-forward computer vision frameworks.

The GSNN and the framework we use for vision problems is completely general. Our next steps will be to apply the GSNN to other vision tasks, such as detection, Visual Question Answering and image captioning. Another inter-

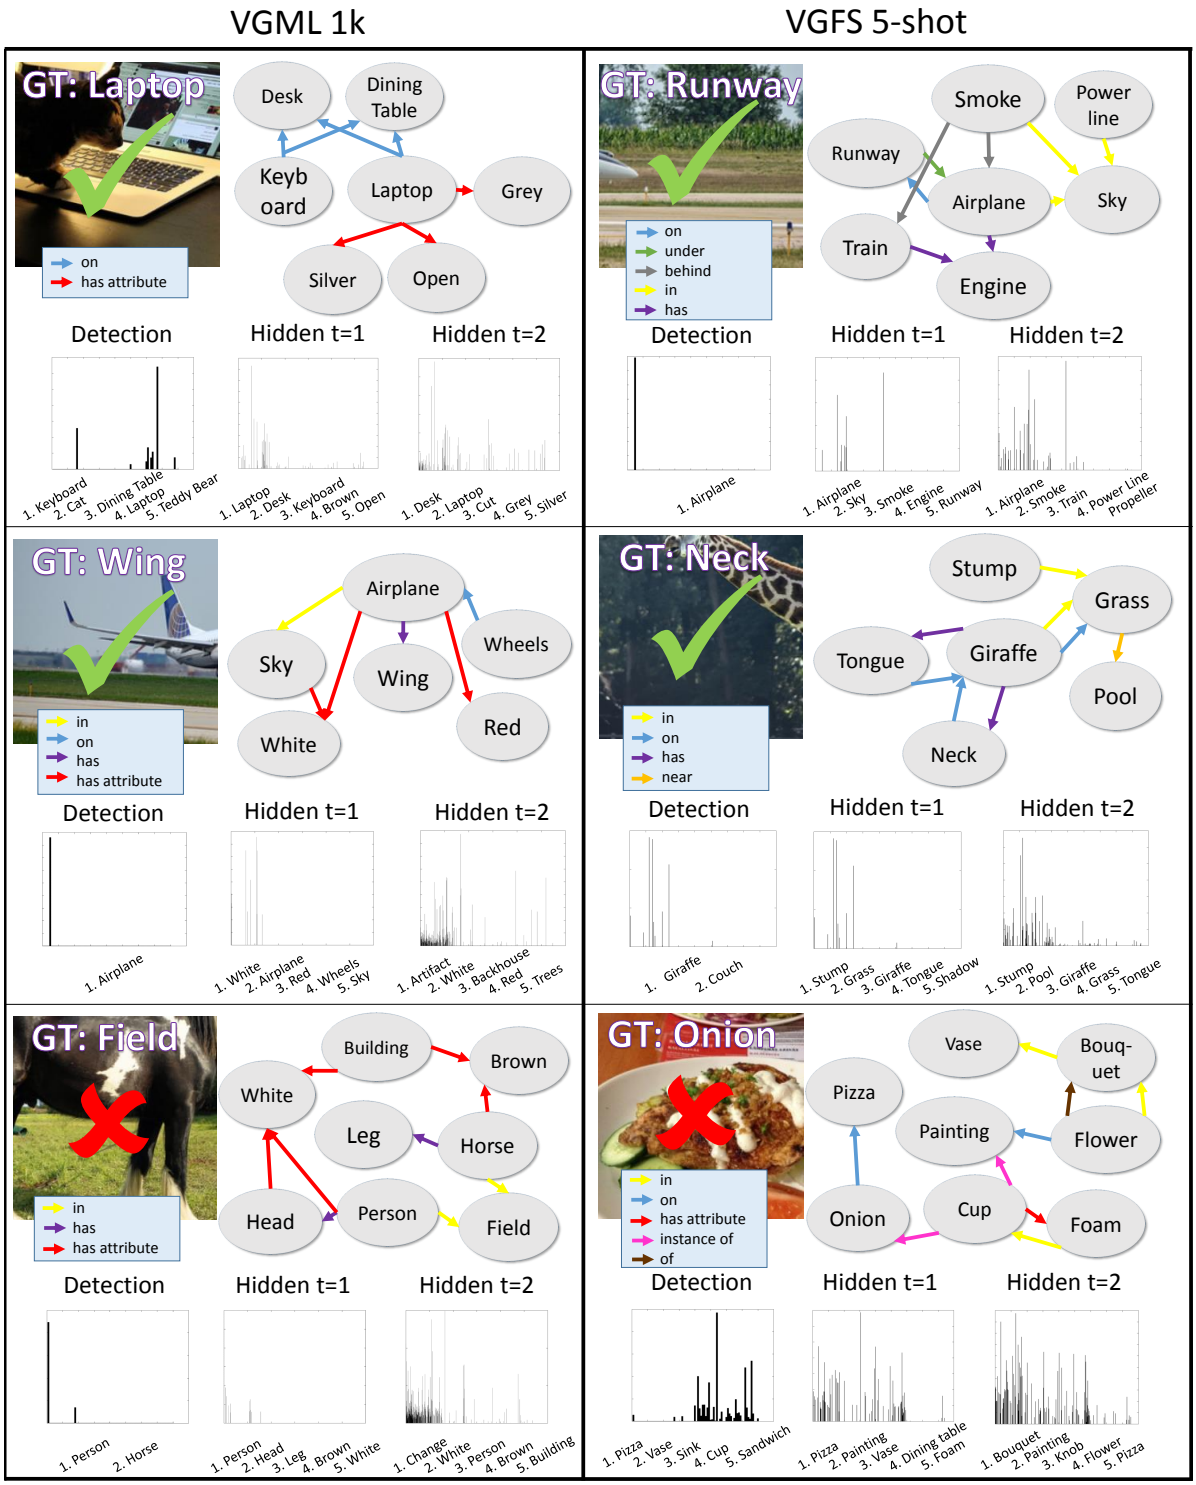

Figure 6. Sensitivity analysis of GSNN in VGML 1k experiment (left) and VGFS 5-shot experiment (right) with the combined graph. Shows part of the expanded knowledge graph and the sensitivity values for each hidden state and detection, showing how important those nodes were for the final classification. The top detections and nodes are printed for convenience. The top and middle rows shows images where the displayed output class was predicted in the final classification while bottom row shows images where it failed to include the class.

esting direction would be to combine the procedure of this work with a system such as NEIL [38] to create a system which builds knowledge graphs and then prunes them to get a more accurate, useful graph for image classification or detection.

Acknowledgements: We would like to thank everyone who took time to review this work and provide helpful comments. This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied of ODNI, IARPA or the US government. The US Government is authorized to reproduce and distribute the reprints for governmental purposed notwithstanding any copyright annotation therein. This material is based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1252522 and ONR MURI N000141612007.

References

[1] K. M. Borgwardt and H.-P. Kriegel. Shortest-path kernels on graphs. *ICDM*, 2005.

[2] J. B. T. Brenden M. Lake, Ruslan Salakhutdinov. One-shot learning by inverting a compositional causal process. *NIPS*, 2013.

[3] J. Bruna, W. Zaremba, A. Szlam, and Y. LeCun. Spectral networks and locally connected networks on graphs. *arXiv preprint arXiv:1312.6203*, 2013.

[4] A. Carlson, J. Betteridge, B. Kisiel, B. Settles, E. R. Hruschka, and T. M. Mitchell. Toward an architecture for never-ending language learning. *AAAI*, 2010.

[5] K. Duan, D. Parikh, D. Crandall, and K. Grauman. Discovering localized attributes for fine-grained recognition. *CVPR*, 2012.

[6] D. K. Duvenaud, D. Maclaurin, J. Iparraguirre, R. Bombarell, T. Hirzel, A. Aspuru-Guzik, and R. P. Adams. Convolutional networks on graphs for learning molecular fingerprints. *NIPS*, 2015.

[7] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html.

[8] A. Farhadi, I. Endres, D. Hoiem, and D. Forsyth. Describing objects by their attributes. *CVPR*, 2009.

[9] M. Gori, G. Monfardini, and F. Scarselli. A new model for learning in graph domains. *IEEE International Joint Conference on Neural Networks*, 2, 2005.

[10] K. Guu, J. Miller, and P. Liang. Traversing knowledge graphs in vector space. In *Empirical Methods in Natural Language Processing (EMNLP)*, 2015.

[11] M. Henaff, J. Bruna, and Y. LeCun. Deep convolutional networks on graph-structured data. *arXiv preprint arXiv:1506.05163*, 2015.

[12] J. Johnson, R. Krishna, M. Stark, L.-J. Li, D. A. Shamma, M. S. Bernstein, and L. Fei-Fei. Image retrieval using scene graphs. *CVPR*, 2015.

[13] D. P. Kingma and J. L. Ba. Adam: A method for stochastic optimization. *ICLR*, 2015.

[14] R. I. Kondor and J. Lafferty. Diffusion kernels on graphs and other discrete input spaces. *ICML*, 2, 2002.

[15] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L.-J. Li, D. A. Shamma, M. Bernstein, and L. Fei-Fei. Visual genome: Connecting language and vision using crowdsourced dense image annotations. 2016.

[16] C. H. Lampert, H. Nickisch, and S. Harmeling. Attribute-based classification for zero-shot visual object categorization. *TPAMI*, 2014.

[17] N. Lao, T. Mitchell, and W. W. Cohen. Random walk inference and learning in a large scale knowledge base. *NIPS*, 2011.

[18] Y. Li and R. Zemel. Gated graph sequence neural networks. *ICLR*, 2016.

[19] T. Lin, M. Maire, S. J. Belongie, R. B. Girshick, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: common objects in context. *ECCV*, 2014.

[20] T. Malisiewicz and A. Efros. Beyond categories: The visual memex model for reasoning about object relationships. *NIPS*, 2009.

[21] V. D. Massa, G. Monfardini, L. Sarti, F. Scarselli, M. Maggini, and M. Gori. A comparison between recursive neural networks and graph neural networks. *IEEE International Joint Conference on Neural Network Proceedings*, 2006.

[22] A. Micheli. Neural network for graphs: A contextual constructive approach. *IEEE Transactions on Neural Networks*, 2009.

[23] G. A. Miller. Wordnet: A lexical database for english. *ACM*, 38, 1995.

[24] M. Niepert, M. Ahmed, and K. Kutzkov. Learning convolutional neural networks for graphs. *arXiv preprint arXiv:1605.05273*, 2016.

[25] F. Orsini, P. Frasconi, and L. D. Raedt. Graph invariant kernels. *IJCAI*, 2015.

[26] Pinar, Yanardag, and S. V. N. Vishwanathan. Deep graph kernels. *KDDM*, 2015.

[27] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. *NIPS*, 2015.

[28] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 115(3):211–252, 2015.

[29] F. Sadeghi, S. K. Divvala, and A. Farhadi. Viske: Visual knowledge extraction and question answering by visual verification of relation phrases. *CVPR*, 2015.

[30] F. Scarselli, M. Gori, A. C. Tsoi, and G. Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 2009.

[31] N. Shervashidze, P. Schweitzer, E. J. van Leeuwen, K. Mehlhorn, and K. M. Borgwardt. Weisfeiler-lehman graph kernels. *JMLR*, 2011.

[32] N. Shervashidze, S. V. N. Vishwanathan, T. H. Petri, K. Mehlhorn, and K. M. Borgwardt. Efficient graphlet kernels for large graph comparison. *AISTATS*, 5, 2009.

[33] A. Shrivastava, S. Singh, and A. Gupta. Constrained semi-supervised learning using attributes and comparative attributes. *ECCV*, 2012.

[34] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

[35] S. V. N. Vishwanathan, N. N. Schraudolph, R. Kondor, and K. M. Borgwardt. Graph kernels. *JMLR*, 2010.

[36] Y.-X. Wang and M. Hebert. Learning to learn: Model regression networks for easy small sample learning. *ECCV*, 2016.

[37] A. Wong and A. L. Yuille. One shot learning via compositions of meaningful patches. *CVPR*, 2015.

[38] X. Xhen, A. Shrivastava, and A. Gupta. Neil: Extracting visual knowledge from web data. *CVPR*, 2013.

[39] X. Zhu, D. Anguelov, and D. Ramanan. Capturing long-tail distributions of object subcategories. *CVPR*, 2014.

[40] Y. Zhu, A. Fathi, and L. Fei-Fei. Reasoning about Object Affordances in a Knowledge Base Representation. In *European Conference on Computer Vision*, 2014.

[41] Y. Zhu, C. Zhang, C. R, and L. Fei-Fei. Building a large-scale multimodal knowledge base system for answering visual queries. 2015.

A. Graph Sensitivity Analysis

Figure 7 shows the graph sensitivity analysis for the VGFS 1-shot experiment on the GSNN combined graph model. The top example shows that it uses the bowl, bottle, pizza, spoon and sandwich detections, many of which connect on the graph to counter, allowing it to reason that the scene contains a counter. The middle example shows that the bed connection to headboard as well as the person detection to bed allows it to reason that it is looking at a bed which relates to headboard. In the failed example on the bottom, the GSNN wrongly detects orange, and the other detections person, tie, horse and truck are not correct and don't have any graph connections to flower, so it fails to classify correctly.

Figure 8 shows the graph sensitivity analysis for the VGML full data experiment on the GSNN combined graph model. The top example shows that using the connection between the cat, person and couch detections, it is able to reason that the scene is on a couch or bed and therefore the connection from both of these to pillow allows it to make the right classification. The middle example shows that it is able to use the connection from dining table, to plate, to cake to reinforce the weak cake detection to correctly classify the image. The failed example on the bottom shows that the network fails to use a detection for dining table which causes it to fail to find any graph connection to wooden. It does expand to bench which has wooden, but it fails to propagate all the way to wooden enough to make a correct classification.

Figure 9 shows the graph sensitivity analysis for the VGML 5k experiment on the GSNN combined graph model. The top expample shows that although train is only the second strongest-used detection, it is able to reason using the connections between train and tracks that it should classify the image as tracks. The middle example shows that detecting TV and Laptop reinforce each other on the graph and cause it to make desk more likely and dining table less likely. The bottom example shows a failed classification where the graph expands on the person detector into less useful states such as car, white and hair. There is no obvious connections in the graph such as surfboard to beach that might have been helpful to make this classification.

Figure 10 shows the graph sensitivity analysis for the VGML 500 experiment on the GSNN combined graph model. The top example shows that using the airplane detection and using the "has" connections from airplane to wing, GSNN is able to classify the given example as wing. In the middle example, the elephant detection is important, which has the obvious connection to trunk. It is also sensitive to the person detection, and this leads to some nodes such as visible becoming important, but it is still correctly able to reason that trunk is in the image. In the failed classification on the bottom row, we see that it uses the dining table and chair detections and even weakly uses the mouse and keyboard detections. Unfortunately, the graph does not contain edges that reinforce a connection between mouse and keyboard, so the graph expands irrelevant nodes such as white and black and fails to classify mouse.

B. Datasets and Graphs

B.1. Dataset and splits

Our dataset splits of Visual Genome into Visual Genome Multi-Label and Visual Genome Few-Shot will be made publicaly available.

The classes we use for VGML are:

person, window, tree, building, sky, shirt, wall, sign, ground, grass, water, pole, head, car, light, hand, plate, train, hair, people, leg, trees, clouds, giraffe, fence, bus, ear, door, floor, chair, elephant, pants, eye, road, leaves, plane, snow, zebra, hat, horse, dog, street, wheel, cat, jacket, shadow, clock, boat, bird, cloud, nose, shoe, line, cow, field, handle, tail, sidewalk, sheep, flower, bench, pizza, helmet, umbrella, leaf, shorts, arm, glass, truck, letter, bear, bag, motorcycle, face, windows, kite, food, bottle, rock, bowl, tile, player, skateboard, post, tire, surfboard, logo, mirror, number, stripe, glasses, bed, roof, picture, flowers, box, shoes, banana, pillow, foot, cap, toilet, tracks, dirt, house, jeans, background, cup, mouth, shelf, laptop, vase, trunk, spot, lights, beach, legs, board, cake, sink, plant, sand, counter, top, airplane, wave, bush, lamp, child, button, flag, paper, desk, writing, brick, seat, neck, wing, coat, lines, book, vehicle, tower, reflection, part, letters, branch, sunglasses, edge, animal, tie, mountain, ocean, lady, rocks, frisbee, cabinet, hill, headlight, keyboard, container, ceiling, skier, eyes, orange, wheels, towel, suitcase, stripes, donut, frame, couch, windshield, broccoli, sandwich, fruit, track, fork, pot, basket, apple, knife, hands, finger, back, bicycle, traffic light, fire hydrant, stop sign, parking meter, backpack, handbag, skis, snowboard, sports ball, baseball bat, baseball glove, tennis racket, wine glass, spoon, carrot, hot dog, potted plant, dining table, tv, mouse, remote, cell phone, microwave, oven, toaster, refrigerator, scissors, teddy bear, hair drier, toothbrush, white, black, blue, green, red, brown, yellow, visible, small, large, wooden, gray, silver, orange, grey, metal, tall, pink, long, standing, dark, here, clear, round, sitting, tan, parked, walking, purple, open, big, there, wood, striped, short, glass, plastic, young, brick, seen, cloudy, gold, this, old, hanging, empty, flying, on, wet, shown, bright, smiling, bare, colorful, playing, concrete, little, beige, closed, dirty, blonde, light, stone, looking, shiny, square, eating, sliced, surfing, skiing, calm, painted, thin, thick, left, dry, near, leather, tiled, cooked, two, close, right, rectangular, grassy, watching, wearing, outside, present, front, grazing, distant, leafy, paved, lit, dead, running, cement, stacked, light brown.

The classes we use for VGFS are:

snow, shadow, cloud, bush, mountain, bread, blanket, cheese, rope, belt, pipe, stick, pillar, river, awning, onion, sun, painting, log, pepperoni, chimney, cockpit, stool, ladder, salad, egg, rice, potato, leash, rose, soap, olive, sausage, tattoo, ketchup, saddle, foam, blender, saucer, hose, bacon, billboard, cage, moss, headband, wristband, newspaper, stroller, puddle, card, power line, mattress, pool, pickle, stairs, buoy, airport, ham, bandana, stump, steering wheel, chandelier, shed, air conditioner, bouquet, overpass, badge, sidewalk, sky, grass, road, flower, arm, face, counter, wave, brick, neck, reflection, windshield, hand, engine, knob, paw, drawer, fur, beak, wetsuit, gravel, tusk, runway, beard, teeth, city, frosting, propeller, smoke, crowd, headboard, tongue.

B.2. Graphs

The graphs we built for our experiments will also be publicly released. Below is some basic information about the graphs we use.

Visual Genome VGML Graph
Num Nodes: 316
Num Edges: 1255
Num Edge Types: 26

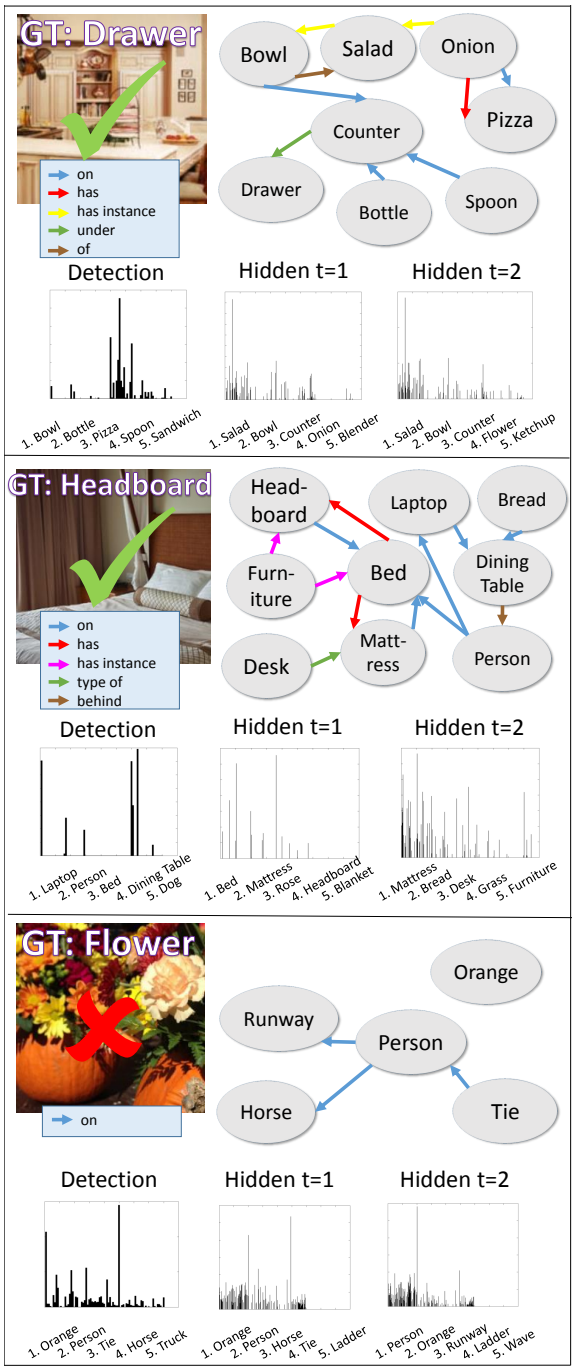

Figure 7. Sensitivity analysis of GSNN in VGFS 1-shot experiment. Shows part of the expanded knowledge graph and the sensitivity values for each hidden state and detection, showing how important those nodes were for the final classification. The top detections and nodes are printed for convenience. The top and middle rows shows images where the displayed output class was predicted in the final classification while bottom row shows images where it failed to include the class.

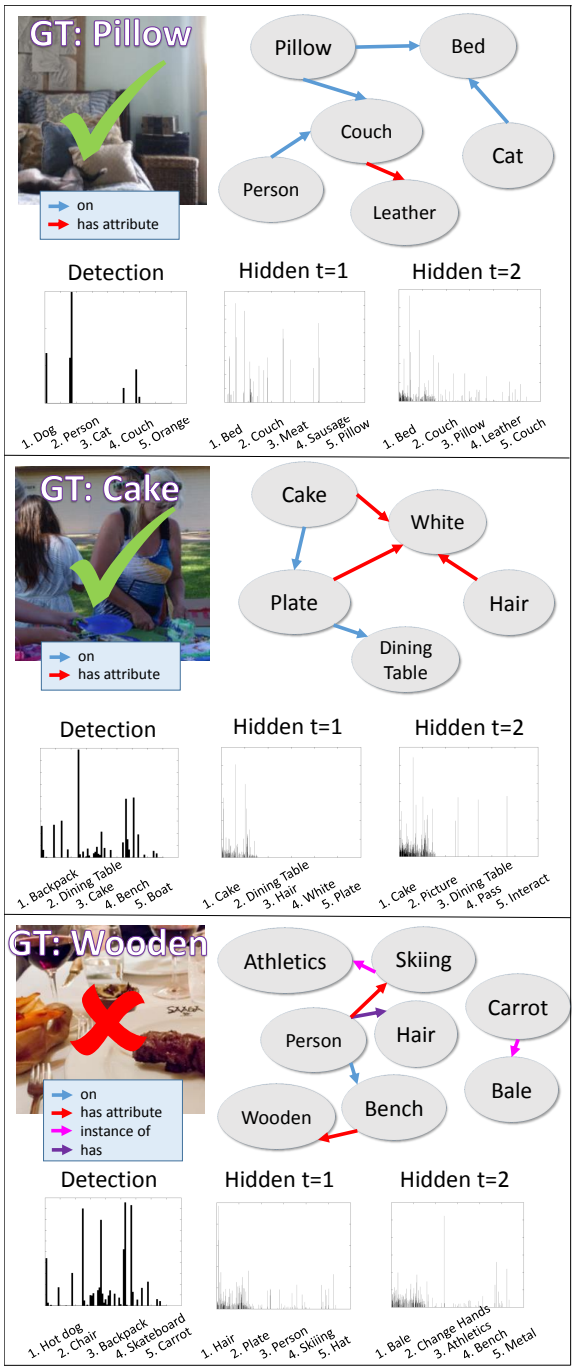

Figure 8. Sensitivity analysis of GSNN in VGML full data experiment. Shows part of the expanded knowledge graph and the sensitivity values for each hidden state and detection, showing how important those nodes were for the final classification. The top detections and nodes are printed for convenience. The top and middle rows shows images where the displayed output class was predicted in the final classification while bottom row shows images where it failed to include the class.

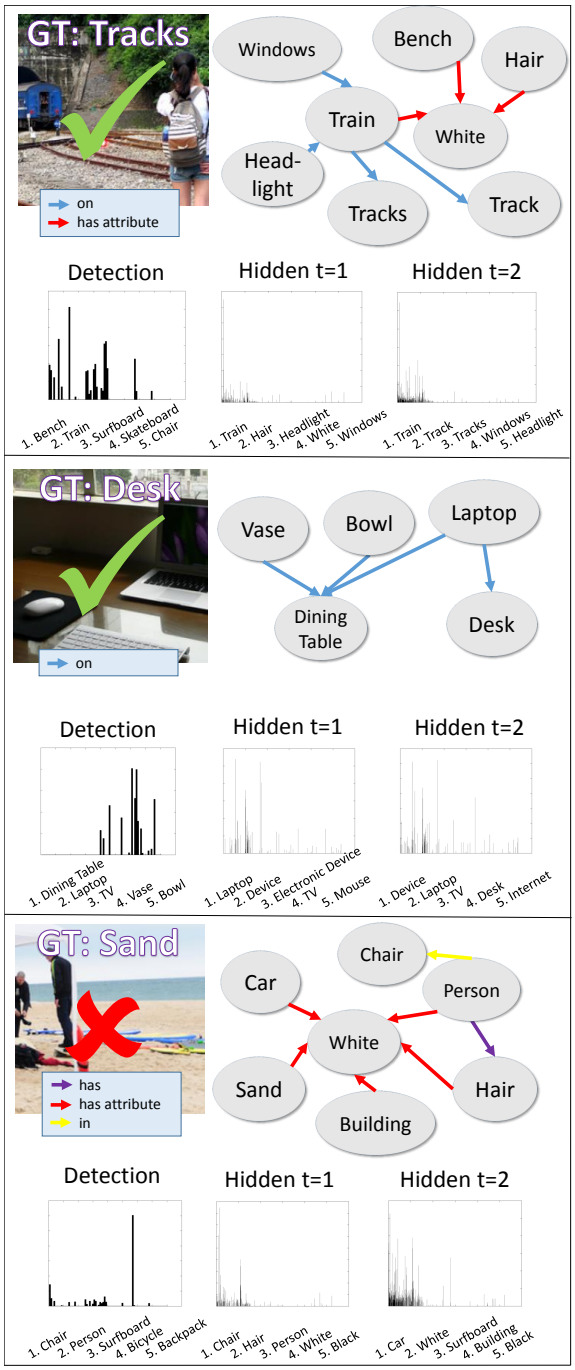

Figure 9. Sensitivity analysis of GSNN in VGML 5k experiment. Shows part of the expanded knowledge graph and the sensitivity values for each hidden state and detection, showing how important those nodes were for the final classification. The top detections and nodes are printed for convenience. The top and middle rows shows images where the displayed output class was predicted in the final classification while bottom row shows images where it failed to include the class.

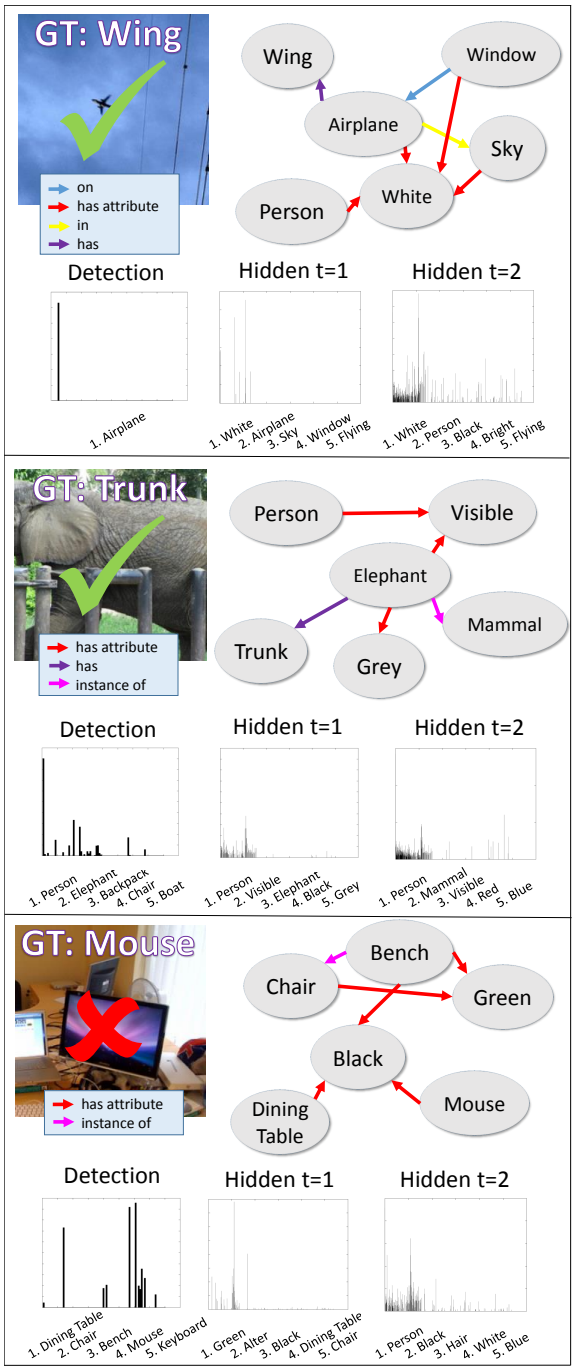

Figure 10. Sensitivity analysis of GSNN in VGML 500 experiment. Shows part of the expanded knowledge graph and the sensitivity values for each hidden state and detection, showing how important those nodes were for the final classification. The top detections and nodes are printed for convenience. The top and middle rows shows images where the displayed output class was predicted in the final classification while bottom row shows images where it failed to include the class.

Combined Visual Genome Wordnet VGML Graph
Num Nodes: 1101
Num Edges: 7666
Num Edge Types: 31

Visual Genome VGFS Graph
Num Nodes: 180
Num Edges: 1289
Num Edge Types: 26

Combined Visual Genome Wordnet VGFS Graph
Num Nodes: 1829
Num Edges: 13080
Num Edge Types: 31

C. Category AP Comparison Figures

Figures 11, 12, 13 and 14 show the category comparison analysis for the VGML full, 5k, 500 and the VGFS 1-shot experiments respectively.

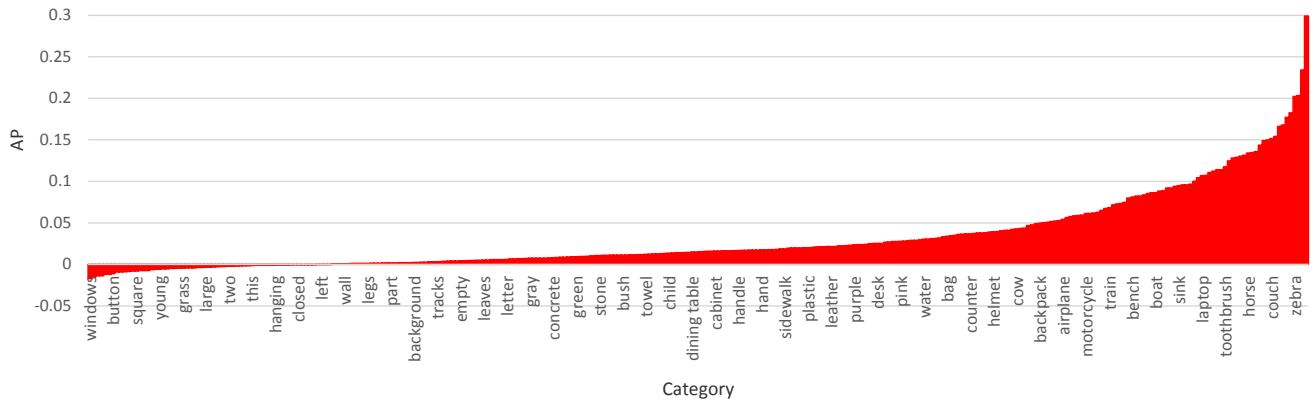

Figure 11. Difference in Average Precision for each of the 316 labels in VGML between our GSNN combined graph model and detection baseline for full data experiment. Top categories: parking meter, cake, frisbee, zebra, kite. Bottom categories: windows, wooden, head, spot, door.

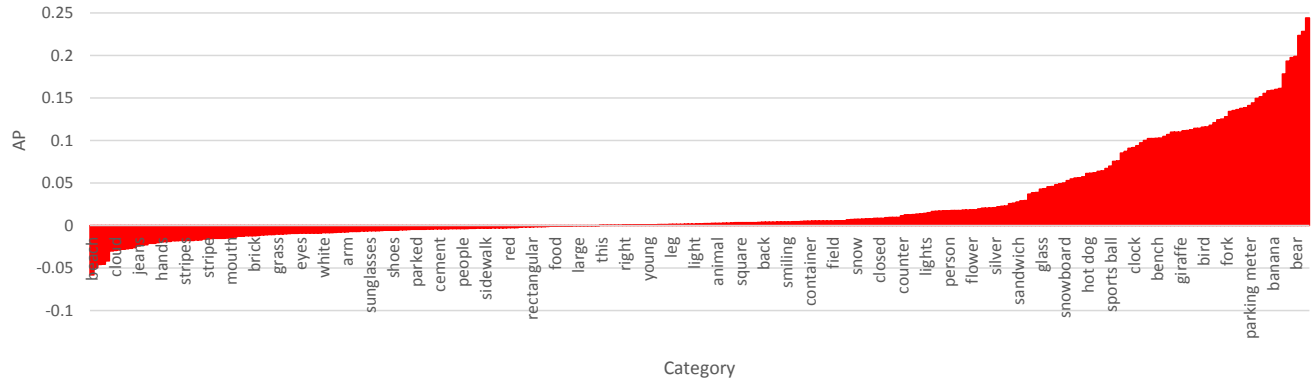

Figure 12. Difference in Average Precision for each of the 316 labels in VGML between our GSNN combined graph model and detection baseline for 5k experiment. Top categories: horse, teddy bear, tv, bear, kite. Bottom categories: beach, ocean, mountain, toothbrush, sand.

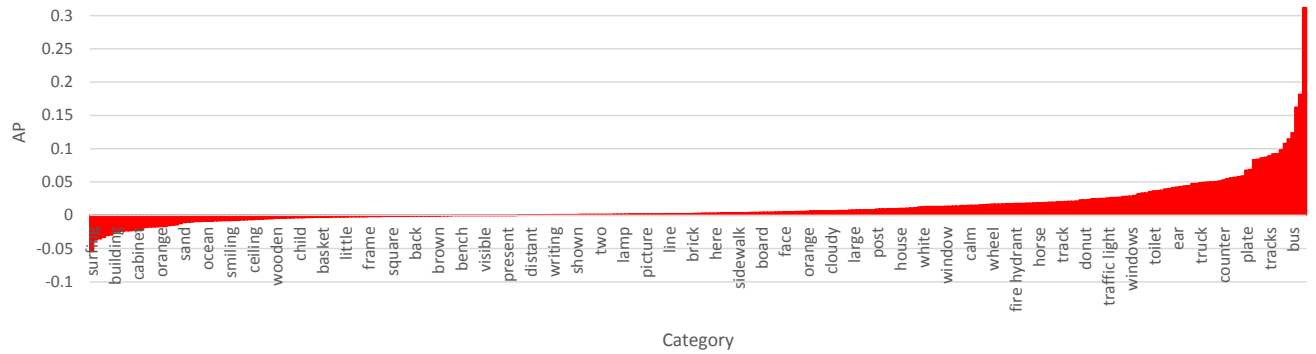

Figure 13. Difference in Average Precision for each of the 316 labels in VGML between our GSNN combined graph model and detection baseline for 500 experiment. Top categories: train, elephant, plane, bus, banana. Bottom categories: surfing, beach, mouse, hair, keyboard.

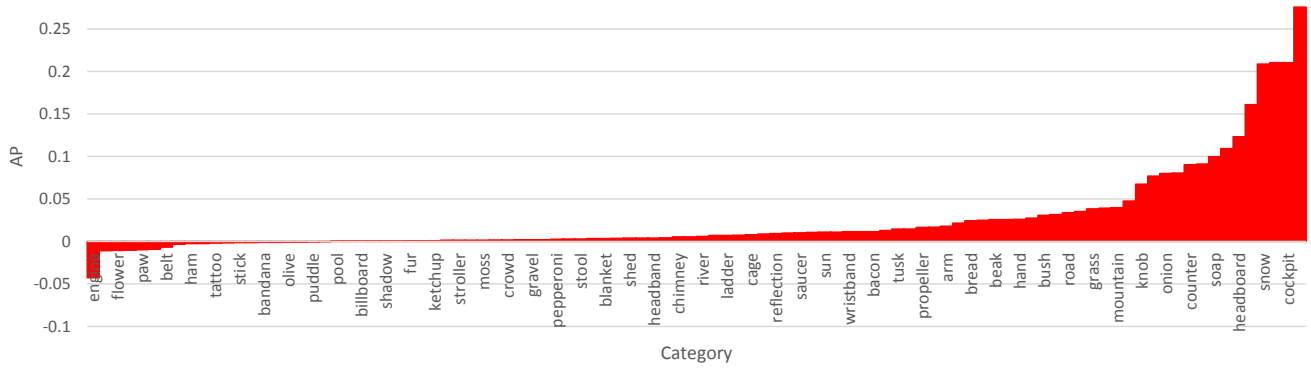

Figure 14. Difference in Average Precision for each of the 100 labels in VGFS between our GSNN combined graph model and detection baseline for 1-shot experiment. Top categories: beak, tusk, runway, snow, headboard. Bottom categories: road, onion, wave, bacon, sidewalk.